Split Sample Empirical Likelihood

Adam Jaeger[1,2] and Nicole Lazar[3]

[1] *Department of Statistical Sciences, Duke University, Durham, NC*

[2] *SAMSI, Research Triangle Park, NC*

[3] *Department of Statistics, University of Georgia, Athens, GA*

arXiv:1703.03312v1 [stat.ME] 9 Mar 2017

Key words and phrases: Parallel computation; Inference; Non-parametric; Parameter estimation

Abstract: We propose a new approach that combines multiple non-parametric likelihood-type components to build a data-driven approximation of the true likelihood function. Our approach is built on empirical likelihood, a non-parametric approximation of the likelihood function. We show the asymptotic behaviors of our approach are identical to those seen in empirical likelihood. We demonstrate that our method performs comparably to empirical likelihood while significantly decreasing computational time.

1 INTRODUCTION

Modern data poses questions of interest that necessitate complex models, atypical probabilistic distributions, and generally contain a very large number of observations and variables. As a result the standard likelihood methodology cannot be properly used owing to the inability to properly express the probability function. Empirical likelihood (Owen, 1988; 1990) is a data driven likelihood that does not require any specification of the probability function. Under mild regularity conditions empirical likelihoods inherit the asymptotic properties of the Fisher likelihood (Wilks, 1938; Qin and Lawless, 1994), and with very few exception (Lazar and Mykland, 1999) permit a Bartlett correction (DiCiccio et al., 1991). This result extends empirical likelihood methods to quantiles and Huber's location M estimate (Huber, 1964). Any functional that has a Fréchet derivative also exhibits the same asymptotic behavior as seen using M estimators, resulting in a large range of applicability. The general form and first-order asymptotics of the empirical likelihood are explored by Qin & Lawless (1994).

Empirical likelihood methods have been extended to many problems such as bivariate means (Owen, 1990), constrained empirical likelihood which includes the creation of a conditional empirical

likelihood, Euclidian likelihood which allows the confidence region to extend beyond the convex hull, and triangular array empirical likelihood which relaxes the assumption of the data being identically distributed (Owen, 1991), regression models, correlation models, ANOVA and variance modeling (Owen, 1991; 2001), the entire class of projection pursuit models (Owen, 1992; Kolaczyk, 1994), time series modeling (Owen, 2001; Kitamura, 1997), incorporating information from multiple moments for a parameter, two sample problems with a common mean, probability measure, incomplete information problems (Qin and Lawless, 1994), ratios of parameters and logistic regression (Qin and Lawless, 1995), partially linear models (Shi and Lau, 2000), missing response problems (Qin and Zhang, 2007), finite population inference (Chen and Qin, 1993), and Bayesian settings (Lazar, 2003; Grendar and Judge, 2010).

Despite all the advantages of empirical likelihood, there are serious computational obstacles that have prevented this method from being commonly used. The empirical likelihood does not result in a closed form; as a consequence computation requires numerical optimization, resulting in non-trivial and time consuming computations. Even if the time consideration is ignored, there are also memory limitations. As sample sizes (and number of parameters) increase the memory requirements for optimization routines can quickly exceed the system memory, making the parameter estimates (let alone any inferential bounds) impossible to compute.

If the computational and memory limitations of empirical likelihood could be solved, this would reintroduce a very flexible and robust method as a potential solution to many modern statistical and scientific questions of interest. For that purpose we introduce a variation of empirical likelihood which maintains all the properties of empirical likelihood and can use parallel computation to dramatically reduce necessary computation time. We call this construct split sample empirical likelihood (SEL). We will show that despite the separation of the data the asymptotic properties of this approach remain consistent with empirical likelihood, resulting in the same asymptotic properties of the parameter estimates and test statistics seen with empirical likelihood.

We start with a review of empirical likelihood. Let $\mathbf{x}_1, \ldots, \mathbf{x}_n$ be d-variate independent identically distributed observations from some cumulative distribution F_0. The empirical likelihood function is

$$L(F) = \prod_{i=1}^{n} dF(\mathbf{x}_i) = \prod_{i=1}^{n} Pr(\mathbf{X} = \mathbf{x}_i) = \prod_{i=1}^{n} u_i.$$

The empirical likelihood function is maximized by the empirical distribution function

$$L(F_n) = \prod_{i=1}^{n} n^{-1},$$

so the empirical likelihood ratio function $R(F) = L(F)/L(F_n)$ can be written as

$$R(F) = \prod_{i=1}^{n} nu_i.$$

Suppose now we are interested in the estimation of a $p \times 1$ parameter $\boldsymbol{\theta}$. We add additional constraints in the form of $r \geq p$ unbiased estimating equations $g(\mathbf{x}, \boldsymbol{\theta})$. These estimating equations generally follow traditional statistics (such as sample means) or the maximum of an optimizing function. We assume the following condition of the estimating equation holds:

$$Eg(\mathbf{X}, \boldsymbol{\theta}_0) = 0,$$

which if we compare to assumptions with normal likelihood equates to

$$E\frac{\partial \log \mathcal{L}(\theta)}{\partial \theta} = 0$$

where $\mathcal{L}(\theta)$ is the parametric likelihood function.

The profile empirical likelihood ratio function is

$$R_E(\boldsymbol{\theta}) = \sup_{\mathbf{u}} \left(\prod_{i=1}^{n} nu_i \mid u_i \geq 0, \sum_{i=1}^{n} u_i = 1, \sum_{i=1}^{n} u_i g(\mathbf{x}_i, \boldsymbol{\theta}) = 0 \right). \tag{1}$$

Provided that $\boldsymbol{\theta}$ is inside the convex hull of the set $g(\mathbf{x}_1, \boldsymbol{\theta}), \ldots, g(\mathbf{x}_n, \boldsymbol{\theta})$ a unique value of Equation 1 exists (Owen, 1988). By definition $R_E(\boldsymbol{\theta}) = 0$ for all $\boldsymbol{\theta}$ not inside the convex hull. The convex hull of the point cloud S is the set

$$\text{Conv}(S) = \left\{ \sum_{i=1}^{|S|} \alpha_i x_i \,\middle|\, (\forall i :\geq 0) \wedge \sum_{i=1}^{|S|} \alpha_i = 1 \right\}.$$

The implications of solutions only existing inside the convex hull will be further explored in Section 2 and Section 4.

The rest of the article is as follows. We formally define the construction of the split sample empirical likelihood in Section 2 along with necessary conditions and notation. Section 3 includes all relevant lemmas, theorems and corollaries which show that SEL has the same asymptotic behavior and forms as empirical likelihood. Section 4 contains a simulation study to demonstrate how

to properly construct a separable empirical likelihood, along with the results indicating the SEL has a significantly shorted computation time and does not significantly deviate from the solutions computed using empirical likelihood. We end with Section 5 which will briefly discuss some practical considerations with our method.

2 SEPARABLE EMPIRICAL LIKELIHOOD

To give an intuitive summary we take our sample and separate into J pieces; with each observation (and all corresponding variables) appearing in only one subset. Although from a purely theoretical standpoint how the original sample is split is unimportant, for practice the sample must be separated so that each subset are fairly uniform (which we will demonstrate in Section 4). Furthermore we are interested in a parameter set which is the same between all data subsets. The next paragraph formally defines the approach.

Let $\mathbf{x}$ be a d dimensional random sample of size n drawn from some cumulative distribution function F_X. We divide $\mathbf{x}$ into J groups, denoting the data of each group $\mathbf{x}^{(j)}$ for $j = 1, \ldots, J$, such that $\cup_{j=1}^{J} \mathbf{x}^{(j)} = \mathbf{x}$ and $\mathbf{x}^{(j)} \cap \mathbf{x}^{(k)} = \emptyset$ for all $j \neq k$. Denote the sample size of each group n_j; note $\sum_{j=1}^{J} n_j = n$. Finally the estimating equations g are the same for all j.

For each subset j the component empirical likelihood is

$$L_E^{(j)}(\boldsymbol{\theta}) = \left(\sup_{\mathbf{u}^{(j)}} \prod_{i=1}^{n_j} u_i^{(j)} \left| \sum_{i=1}^{n_j} u_i^{(j)} g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}) = 0,\ u_i^{(j)} \geq 0,\ \sum_{i=1}^{n_j} u_i^{(j)} = 1 \right. \right),$$

and the split sample empirical likelihood (SEL) function is

$$L_S(\boldsymbol{\theta}) = \prod_{j=1}^{J} L_E^{(j)}(\boldsymbol{\theta})..$$ (2)

Note that the component empirical likelihood is simply the empirical likelihood consisting of a subset of the data.

We use the following notation

$$\ell_E^{(j)}(\boldsymbol{\theta}) = -\log(L_E^{(j)}(\boldsymbol{\theta}))$$

to denote the negative log of the component empirical likelihood, which leads to the negative log split sample empirical likelihood function

$$\ell_S(\boldsymbol{\theta}) = -\log\left(\prod_{j=1}^{J} L_E^{(j)}(\boldsymbol{\theta}) \right) = \sum_{j=1}^{J} \ell_E^{(j)}(\boldsymbol{\theta}).$$ (3)

Equation 3 shows the connection between empirical likelihood and SEL; SEL is defined as the log of the product of J component empirical likelihoods.

Exactly like empirical likelihood $\ell_S(\boldsymbol{\theta})$ is only defined inside the convex hull. Let $\mathcal{C}_j$ be the convex hull for each component likelihood. The convex hull of ℓ_S is

$$\mathcal{C}_S = \cap_{j=1}^{J} \mathcal{C}_j.$$

We define $\ell_S(\boldsymbol{\theta}) = 0$ for all values of $\boldsymbol{\theta}$ outside the convex hull.

As previously stated there are no theoretical restrictions on how to separate the sample, but there are practical considerations. The two primary considerations for splitting the data are keeping n_j as equal as possible, i.e. $n_j \approx n/J$ for all j, and forcing the data structure to be as similar as possible in each component. To show a toy example of the second point, if we had univariate data and $J = 2$ we would not put all smaller values in one component and all larger values in the second component. Figure 1 shows an example of this sort of split.

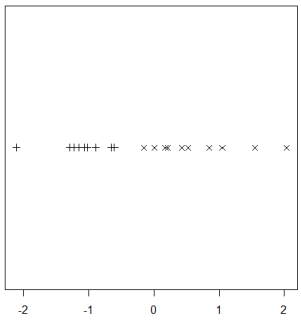


Figure 1: Univariate Data Split into smaller half and larger half. Note there is no overlap between the two sets (+ and x), so the convex hull of the intersection of the two sets is empty.

If the data were split as such, $\ell_S(\boldsymbol{\theta}) = 0$ for all values of θ. For this reason the data need to be separated so that they are representative of the original data, making the convex hull as large as possible. Related to this issue is for a fixed sample size the convex hull will almost always shrink as J increases, and can even result in the empty set. We will show a sorting method in Section 4 designed to minimize the possibility of the convex hull being unreasonably small.

We also need to assume that there exists a unique global maximum of $\ell_S(\boldsymbol{\theta})$, and further more that the maximum does not exist right on the edge of the convex hull in order to generate confidence

intervals/regions and perform hypothesis tests. As such we make the following assumptions.

Assumption 1. Using the notation defined for the component empirical likelihood and the split sample empirical likelihood we will assume the following conditions hold.

1. A unique global maximum of $\ell^{(j)}(\boldsymbol{\theta})$ for all J exists in the convex hull $\mathcal{C}_S$,
2. $|\partial \ell^{(j)}(\boldsymbol{\theta})/\partial \boldsymbol{\theta}| > 0$ exists in the convex hull $\mathcal{C}_S$.

We give proof of why these conditions are sufficient in the Appendix.

Finally, the maximum split sample empirical likelihood (MSEL) estimator is the value of $\boldsymbol{\theta}$ that maximizes $L_S(\boldsymbol{\theta})$; or more formally:

$$\hat{\boldsymbol{\theta}}_S = \sup_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta}) = \inf_{\boldsymbol{\theta}} \ell_S(\boldsymbol{\theta}).$$

For the remainder of the paper we will use the following notation

$$E\left(g(\mathbf{X}, \boldsymbol{\theta}_0) g^{\mathrm{T}}(\mathbf{X}, \boldsymbol{\theta}_0)\right) = E\left(g g^{\mathrm{T}}\right)$$
$$E\left(\left.\frac{\partial g(\mathbf{X}, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}\right) = E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)$$

where $\boldsymbol{\theta}_0$ is the true parameter value.

3 ASYMPTOTIC PROPERTIES

Several of our proofs utilize results shown in (Qin and Lawless, 1994), which we give here and in context of our notation. We will use the prefix QL to denote that these lemma's are slight variations from (Qin and Lawless, 1994).

Lemma (QL1) Assume Egg^{T} is positive definite with rank p, $\partial g(x^{(j)}, \theta)/\partial \theta$ is continuous in a neighborhood of the true value θ_0, $\|\partial g(x^{(j)}, \theta)/\partial \theta\|$ and $\|g(x^{(j)}, \theta)\|^3$ are bounded by some integrable function $G(x)$ in this neighborhood. Then, as $n \to \infty$, with probability 1 $\ell^{(j)}$ for $j = 1, \ldots, J$ attains its minimum value at some point $\hat{\theta}$ in the interior of the ball $\|\theta - \theta_0\| \leq n^{-1/3}$, and $\hat{\theta}$ and $\hat{t} = t(\hat{\theta})$ satisfy

$$Q_{1j}(\hat{\boldsymbol{\theta}}_S, \hat{\mathbf{t}}_j) = \frac{1}{n_j} \sum_{i=1}^{n_j} (1 + \hat{\mathbf{t}}_j^{\mathrm{T}} g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}))^{-1} g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}) = 0$$

for all $j = 1, \ldots, J$.

We do not offer proof of Lemma QL1 as it follows directly from Qin and Lawless (1994, Lemma 1).

We now provide a lemma similar to Qin and Lawless (1994, Lemma 1) in order to show that the MSEL maximizes the SEL for a sufficiently large n, which are also necessary results for deriving the asymptotic distribution of $\hat{\boldsymbol{\theta}}_S$ and the separable empirical likelihood ratio test statistic. We provide a proof for Lemma 1 in the Appendix.

Lemma 1. Under Assumption 1 and the assumptions stated in Lemma QL1 $\ell_S(\boldsymbol{\theta})$ attains its minimum value at some point $\hat{\boldsymbol{\theta}}_S$ in the interior of the ball $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\| \leq n^{-1/3}$ with probability 1 as $n \to \infty$.

Furthermore $\hat{\boldsymbol{\theta}}_S$ and $\hat{\boldsymbol{t}}_j = \mathbf{t}_j(\hat{\boldsymbol{\theta}}_S)$ for all j satisfy

$$Q_2(\hat{\boldsymbol{\theta}}_S, \hat{\boldsymbol{t}}_1, \ldots, \hat{\boldsymbol{t}}_J) = \sum_{j=1}^{J} \left(\frac{1}{n_j} \sum_{i=1}^{n_j} (1 + \mathbf{t}_j^{\mathrm{T}} g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}))^{-1} \left(\frac{\partial g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^{\mathrm{T}} \mathbf{t}_j \right) = 0$$

where $\boldsymbol{t}_j$ is the $r \times 1$ vector given as the solution to

$$\sum_{i=1}^{n_j} \left(1 + \boldsymbol{t}_j^{\mathrm{T}} g(x_i^{(j)}, \boldsymbol{\theta}) \right) g(x_i^{(j)}, \boldsymbol{\theta}) = 0.$$

The next two theorems, Theorem 1 and 2, pertain to the asymptotic behavior of the MSEL and the likelihood ratio test statistic using SEL. We first

Theorem 1. Assume that $\partial^2 g(x^{(j)}, \theta)/\partial\theta\partial\theta^{\mathrm{T}}$ is continuous in θ in a neighborhood of the true value θ_0 for $j = 1, \ldots, J$ along with the Assumption 1 and the Assumptions stated in Lemma QL1. Then if $\|\partial^2 g(x^{(j)}, \theta)/\partial\theta\partial\theta^{\mathrm{T}}\|$ for $j = 1, \ldots, J$ can be bounded by some integrable function $G(x)$ in the neighborhood we have for n sufficiently large

$$\sqrt{n}(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0) \to N(0, W_{\boldsymbol{\theta}}^{-1})$$

where

$$W_{\boldsymbol{\theta}} = E\left(\frac{\partial g}{\partial \boldsymbol{\theta}} \right)^{\mathrm{T}} \left\{ E\left(g g^{\mathrm{T}} \right) \right\}^{-1} E\left(\frac{\partial g}{\partial \boldsymbol{\theta}} \right).$$

Note that the distribution of the maximum empirical likelihood estimator has the same distribution as $\hat{\boldsymbol{\theta}}_S$ (see Qin and Lawless, 1994, Theorem 1), and in the case of $J = 1$ the separable empirical likelihood function is equal to the empirical likelihood function.

Analogously to other likelihood functions, we develop a SEL ratio statistic in order to test hypotheses, and by extension compute confidence regions.

Theorem 2. Let $\boldsymbol{\theta} = [\boldsymbol{\phi}^{\mathrm{T}}, \boldsymbol{\nu}^{\mathrm{T}}]^{\mathrm{T}}$ be a p dimensional vector where $\boldsymbol{\phi}$ is a $q \times 1$ vector and $\boldsymbol{\nu}$ is a $(p - q) \times 1$ vector of nuisance parameters. Under Assumption 1 and the assumptions from QL1

and Lemma 1, the profile split sample empirical likelihood ratio test statistic for $H_0 : \boldsymbol{\phi} = \boldsymbol{\phi}_0$ is

$$T = 2\ell_S(\boldsymbol{\phi}_0, \hat{\boldsymbol{\nu}}_S(\boldsymbol{\phi}_0)) - 2\ell_S(\hat{\boldsymbol{\phi}}_S, \hat{\boldsymbol{\nu}}_S)$$

where $\hat{\boldsymbol{\nu}}_S(\boldsymbol{\phi}_0)$ minimizes $\ell_S(\boldsymbol{\phi}, \boldsymbol{\nu})$ with respect to $\boldsymbol{\phi}_0$. Under H_0

$$T \to \chi^2(q)$$

as $n \to \infty$.

Corollary 1. The SEL ratio statistic for testing $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ is

$$T = 2\ell_S(\boldsymbol{\theta}_0) - 2\ell_S(\hat{\boldsymbol{\theta}}_S),$$

Given the assumptions stated in Theorem 2, under H_0

$$T \to \chi^2(p)$$

as $n \to \infty$.

We provide proofs of Theorems (1) and (2) in the Appendix. The proof of Corollary (1) follows directly from Theorem (2) by taking $\boldsymbol{\nu} = \emptyset$.

4 SIMULATIONS

The simulations are designed to demonstrate how to appropriate divide the data into the J likelihood components, and confirm that the resultant estimates for a given data set do not differ significantly from the estimators obtained using empirical likelihood. The simulations will also show that by optimizing each likelihood component in parallel the computation time is significantly decreased.

The data are generated from a bivariate Poisson which has the following probability mass function

$$P(X = x, Y = y) = e^{-(\theta_1+\theta_2+\theta_0)} \frac{\theta_1^x}{x!} \frac{\theta_2^y}{y!} \sum_{i=0}^{\min(x,y)} \binom{x}{i} \binom{y}{i} i! \left(\frac{\theta_0}{\theta_1\theta_2}\right)^i.$$

We will use $\theta_1 = 2$, and $\theta_2 = \theta_0 = 3$ along with a sample size of 2000. Our focus will be on inference of the means of X and Y along with the correlation between X and Y, which in this instance is $EX = \mu_x = \theta_1 + \theta_0 = 5$, $EY = \mu_y = \theta_2 + \theta_0 = 6$, and $Corr(X, Y) = \rho = \theta_0/(\theta_1\theta_2) = 1/2$. The

estimating equations for each empirical likelihood component are

$$g = \begin{pmatrix} x_i^{(j)} - \mu_x \\ y_i^{(j)} - \mu_y \\ \frac{(x_i^{(j)} - \tilde{\mu}_x)(y_i^{(j)} - \tilde{\mu}_y)}{\sqrt{\tilde{\sigma}_x^2 \tilde{\sigma}_y^2}} - \rho \end{pmatrix}$$

where

$$\tilde{\mu}_x = \sum_{i=1}^{n_j} u_i^{(j)} x_i^{(j)}$$

$$\tilde{\mu}_y = \sum_{i=1}^{n_j} u_i^{(j)} y_i^{(j)}$$

$$\tilde{\sigma}_x^2 = \sum_{i=1}^{n_j} u_i^{(j)} \left(x_i^{(j)} - \tilde{\mu}_x \right)^2$$

$$\tilde{\sigma}_y^2 = \sum_{i=1}^{n_j} u_i^{(j)} \left(y_i^{(j)} - \tilde{\mu}_y \right)^2$$

There are multiple approaches to assigning the data to each component provided the assignment guarantees some level of uniformity among each component. The sort algorithm we used is as follows:

1. Extract the J smallest and J largest sets from the original data. Here we would consider $(1, 1)$ to be a smaller pair than $(0, 2)$ with a similar notion holding for largest.

2. The maximum and minimum are assigned to $j = 1$, the second largest and second smallest to $j = 2$ and so on.

3. The remainder of the data are sorted largest to smallest and a systematic sample approach is used.

As a result the observation pairs are kept intact, and each likelihood component has a representative sample in relation to the original data. This is not the only approach to data assignment; any sort that results in general data uniformity in each likelihood component is recommended. Although we have not examined the scenario where the number of observations are extremely large in relation to J, but our assumption is for a large number of observations and a relatively small J, sorting becomes less critical. The primary purpose of any sorting is to make the convex hull of ℓ_S as a large as possible.

The simulations compare the results obtained using empirical likelihood for the moment estimators described above. Furthermore we look at two sided hypothesis tests of $H_0 : \rho = /1/2$ and

$H_0 : \mu_x = 5$ and $\mu_y = 6$. We examine both deviation from the results obtained with empirical likelihood and include mean computation times for each procedure. Our results are shown below.

[*SIMULATIONS IN PROGRESS*]

5 DISCUSSION

The immediate question that arises is: what is the optimal number of components? This depends on several factors, most noticeably the sample size. As our simulations show there are diminishing returns from a computation time standpoint as J gets larger. From our own experiences we recommend that each empirical likelihood component have at minimum 25 observations if working with a single parameter to ensure that the optimization algorithm will converge for each piece. The number of parameters (including nuisance) and the complexity of the estimating equations increase the necessary minimum number of observations in each likelihood component. Ultimately the goal is to ensure that the convex hull of the split sample empirical likelihood is sufficiently large so that the maximum value and critical values exists.

BIBLIOGRAPHY

Chen, J. and Qin, J. (1993). Empirical likelihood estimation for finite populations and the effective usage of auxiliary information. *Biometrika*, 80(1):107–116.

DiCiccio, T. J., Hall, P., and Romano, J. P. (1991). Empriical likelihood is Bartlett-correctable. *The Annals of Statistics*, 19(2):1053–1061.

Grendar, M. and Judge, G. G. (2010). Revised empirical likelihood. Technical Report 1106, CUDARE Working Paper Series, University of California at Berkeley.

Huber, P. J. (1964). Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35(1):73–101.

Kitamura, Y. (1997). Empirical likelihood methods with weakly dependent processes. *The Annals of Statistics*, 25(5):2084–2102.

Kolaczyk, E. D. (1994). Empirical likelihood for generalized linear models. *Statistica Sinica*, 4(1):199–218.

Lazar, N. A. (2003). Bayesian empirical likelihood. *Biometrika*, 90(2):319–326.

Lazar, N. A. and Mykland, P. A. (1999). Empirical likelihood in the presence of nuisance parameters. *Biometrika*, 86(1):203–211.

Owen, A. B. (1988). Empirical likelihood ratio confidence intervals for a single functional. *Biometrika*, 75(2):237–249.

Owen, A. B. (1990). Empirical likelihood ratio confidence regions. *The Annals of Statistics*, 18(1):90–120.

Owen, A. B. (1991). Empirical likelihood for linear models. *The Annals of Statistics*, 19(4):1725–1747.

Owen, A. B. (1992). Empirical likelihood and generalized projection pursuit. Technical Report 393, Department of Statistics, Stanford University.

Owen, A. B. (2001). *Empirical Likelihood*. Chapman & Hall/CRC, Boca Raton.

Qin, J. and Lawless, J. (1994). Empirical likelihood and general estimating equations. *The Annals of Statistics*, 22(1):300–325.

Qin, J. and Lawless, J. (1995). Estimating equations, empirical likelihood and constraints on parameters. *The Canadian Journal of Statistics*, 23(2):145–159.

Qin, J. and Zhang, B. (2007). Empirical-likelihood-based inference in missing response problems and its application in observational studies. *Journal of the Royal Statistical Society Series B*, 69(1):101–122.

Rao, C. R. (1973). *Linear Statistical Inference and Its Applications*. Wiley, New York.

Shi, J. and Lau, T. (2000). Empirical likelihood for partially linear models. *Journal of Multivariate Analysis*, 72(1):132–148.

Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62.

APPENDIX

Proof of Assumption 1. Let $f_1, \ldots, f_J$ be functions, each with solutions within a convex hull C_j. Define the intersection of the convex hulls $C_S = \cap_{j=1}^{J} \mathcal{C}_j$. Assume that the maximum value for all f_j exists and C_S and that there exists non zero derivatives of all functions some $\varepsilon \neq 0$ from the maximum. Given these two assumptions each function is concave in C_S. We then have that the sum of concave function is also concave, i.e. $f = \sum_{j=1}^{J} f_j$ is also concave in C_S. Therefore there exists a unique maximum of f in C_S. □

Proof of Theorem 1. As shown in Qin and Lawless (1994) by Taylor expansion of $Q_{1j}(\hat{\boldsymbol{\theta}}_S, \hat{\mathbf{t}}_j)$ around $\boldsymbol{\theta}_0$ and 0

$$\begin{aligned} 0 &= Q_{1j}(\hat{\boldsymbol{\theta}}_S, \hat{\mathbf{t}}_j) \\ &= Q_{1j}(\boldsymbol{\theta}_0, 0) + \frac{\partial}{\partial \boldsymbol{\theta}} Q_{1j}(\boldsymbol{\theta}_0, 0)(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0) + \frac{\partial}{\partial \mathbf{t}_j^{\mathrm{T}}} Q_{1j}(\boldsymbol{\theta}_0, 0)(\hat{\mathbf{t}}_j - 0) + o_p(\delta_j) \end{aligned}$$

where $\delta_j = \|\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0\| + \|\hat{\mathbf{t}}_j\|$.

By Taylor expansion of $Q_2(\hat{\boldsymbol{\theta}}_S, \hat{\mathbf{t}}_1, \ldots, \hat{\mathbf{t}}_J)$ around $\boldsymbol{\theta}_0, 0, \ldots, 0$

$$\begin{aligned} 0 &= Q_2(\hat{\boldsymbol{\theta}}_S, \hat{\mathbf{t}}_1, \ldots, \hat{\mathbf{t}}_J) \\ &= Q_2(\boldsymbol{\theta}_0, 0, \ldots, 0) + \frac{\partial}{\partial \boldsymbol{\theta}} Q_2(\boldsymbol{\theta}_0, 0, \ldots, 0)(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0) \\ &\quad + \sum_{j=1}^{J} \frac{\partial}{\partial \mathbf{t}_j^{\mathrm{T}}} Q_2(\boldsymbol{\theta}_0, 0, \ldots, 0)(\hat{t}_j - 0) + o_p(\delta) \end{aligned}$$

where $\delta = \|\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0\| + \sum \|\hat{\mathbf{t}}_j\|$. From here we can show

$$\begin{aligned} 0 = \sum_{j=1}^{J} \Bigg(& \frac{\partial}{\partial \mathbf{t}_j^{\mathrm{T}}} Q_2(\boldsymbol{\theta}_0, 0, \ldots, 0) \left(-\frac{\partial Q_{1j}(\boldsymbol{\theta}_0, 0)}{\partial \mathbf{t}_j^{\mathrm{T}}} \right)^{-1} \times \\ & \left(Q_{1j}(\boldsymbol{\theta}_0, 0) + \frac{\partial}{\partial \boldsymbol{\theta}} Q_{1j}(\boldsymbol{\theta}_0, 0)(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0) \right) \Bigg) + o_p(\max(\delta_j, \delta)). \end{aligned}$$

The derivatives of Q_{1j} and Q_2 with respect to $\boldsymbol{\theta}$ and $\mathbf{t}_j$ are

$$\begin{aligned} \frac{\partial Q_{1j}(\boldsymbol{\theta}_0, 0)}{\partial \boldsymbol{\theta}} &= \frac{1}{n_j} \sum_{i=1}^{n_j} \frac{\partial g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}} \to E\left(\frac{\partial g}{\partial \boldsymbol{\theta}} \right) \\ \frac{\partial Q_{1j}(\boldsymbol{\theta}_0, 0)}{\partial \mathbf{t}_j^{\mathrm{T}}} &= -\frac{1}{n_j} \sum_{i=1}^{n_j} g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}_0) g_j^{\mathrm{T}}(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}_0) \to -E\left(g g^{\mathrm{T}} \right) \\ \frac{\partial Q_2(\boldsymbol{\theta}_0, 0, \ldots, 0)}{\partial \mathbf{t}_j^{\mathrm{T}}} &= \frac{1}{n_j} \sum_{i=1}^{n_j} \left(\frac{\partial g(x_i^{(j)}, \boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}} \right)^{\mathrm{T}} \to E\left(\frac{\partial g}{\partial \boldsymbol{\theta}} \right)^{\mathrm{T}} \end{aligned}$$

so

$$\begin{aligned} 0 = &\sum_{j=1}^{J} \left(E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)^{\mathrm{T}} \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1} \left(-Q_{1j}(\boldsymbol{\theta}_0, 0)\right) \right) \\ &- \sum_{j=1}^{J} \left(E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)^{\mathrm{T}} \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1} E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right) \left(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0\right) \right) \\ &+ o_p(1) \qquad \Rightarrow \end{aligned}$$

$$\left(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0\right) = J^{-1} W_{\boldsymbol{\theta}}^{-1} E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)^{\mathrm{T}} \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2} \sum_{j=1}^{J} \left(- \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2} Q_{1j}(\boldsymbol{\theta}_0, 0) \right) + o_p(1)$$

where

$$W_{\boldsymbol{\theta}} = E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)^{\mathrm{T}} \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1} E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right).$$

Recall

$$Q_{1j}(\boldsymbol{\theta}_0, 0) = \frac{1}{n_j} \sum_{i=1}^{n_j} g(\mathbf{x}_i^{(j)}, \boldsymbol{\theta}_0)$$

so $\sqrt{n_j} Q_{1j}(\boldsymbol{\theta}_0, 0)$ converges to a normal distribution with $E\left(\sqrt{n_j} Q_{1j}(\boldsymbol{\theta}_0, 0)\right) = \boldsymbol{\theta}_0$. Furthermore we have $Q_{1j}(\boldsymbol{\theta}_0, 0) \perp\!\!\!\perp Q_{1k}(\boldsymbol{\theta}_0, 0)$ for all $j \neq k$. If we assume that n_j is approximately equal for each j and increases at the same rate we have as n (the total sample size) gets sufficiently large

$$-\sqrt{n} \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2} Q_{1j}(\boldsymbol{\theta}_0, 0) = -\sqrt{n/J} \left\{JE\left(gg^{\mathrm{T}}\right)\right\}^{-1/2} Q_{1j}(\boldsymbol{\theta}_0, 0) \to N(0, J)$$

so

$$\sum_{j=1}^{J} -\sqrt{n/J} \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2} Q_{1j}(\boldsymbol{\theta}_0, 0) \to N(0, J^2).$$

Finally

$$\begin{aligned} E\left(\sqrt{n}(\hat{\boldsymbol{\theta}}_S - \boldsymbol{\theta}_0)\right)^2 &= J^{-1} W_{\boldsymbol{\theta}}^{-1} E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)^{\mathrm{T}} \left\{E(gg^{\mathrm{T}})\right\}^{-1/2} J^2 \left\{E(gg^{\mathrm{T}})\right\}^{-1/2} E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right) W_{\boldsymbol{\theta}}^{-1} J^{-1} \\ &= W_{\boldsymbol{\theta}}^{-1} \end{aligned}$$

which completes the proof.

□

Proof of Theorem 2. Using the notation and results established in Lemma (1) and Theorem (1) the test statistic can be expressed as

$$T = \sum_{j=1}^{J}\sum_{k=1}^{J}\Bigg(\sqrt{n_j}Q_{1j}^{\mathrm{T}}(\boldsymbol{\theta}_0,0)\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1} \times\left\{E\left(\frac{\partial g}{\partial\boldsymbol{\theta}}\right)J^{-1}W_{\boldsymbol{\theta}}^{-1}E\left(\frac{\partial g}{\partial\boldsymbol{\theta}}\right)^{\mathrm{T}}-E\left(\frac{\partial g}{\partial\boldsymbol{\nu}}\right)J^{-1}W_{\boldsymbol{\nu}}^{-1}E\left(\frac{\partial g}{\partial\boldsymbol{\nu}}\right)^{\mathrm{T}}\right\}\times \left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1}\sqrt{n_k}Q_{1k}(\boldsymbol{\theta}_0,0)\Bigg)+o_p(1)$$

where

$$W_{\boldsymbol{\nu}} = E\left(\left.\frac{\partial g(\mathbf{X},\boldsymbol{\theta})}{\partial\boldsymbol{\nu}}\right|_{\boldsymbol{\nu}=\boldsymbol{\nu}_0}\right)^{\mathrm{T}}\left\{Egg^{\mathrm{T}}\right\}^{-1}E\left(\left.\frac{\partial g(\mathbf{X},\boldsymbol{\theta})}{\partial\boldsymbol{\nu}}\right|_{\boldsymbol{\nu}=\boldsymbol{\nu}_0}\right).$$

and $W_{\boldsymbol{\theta}}$ is as defined in Theorem (1). Assuming that all n_j are increasing at approximately the same rate we have for n sufficiently large $\sqrt{n_j}\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}Q_{1j}(\boldsymbol{\theta}_0,0)$ are asymptotically independent standard normal for all j, so we can rewrite the test statistic as

$$\begin{bmatrix} Z_1 & \dots & Z_J \end{bmatrix} J^{-1}A \begin{bmatrix} Z_1 \\ \dots \\ Z_J \end{bmatrix}$$

where each block of A is

$$\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}\left\{E\left(\frac{\partial g}{\partial\boldsymbol{\theta}}\right)W_{\boldsymbol{\theta}}^{-1}E\left(\frac{\partial g}{\partial\boldsymbol{\theta}}\right)^{\mathrm{T}}-E\left(\frac{\partial g}{\partial\boldsymbol{\nu}}\right)W_{\boldsymbol{\nu}}^{-1}E\left(\frac{\partial g}{\partial\boldsymbol{\nu}}\right)^{\mathrm{T}}\right\}\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}$$

and $Z_1,\dots,Z_J$ are independent standard normal variables.

To complete the proof we need the following proposition.

Proposition 1. Let Z be from a multivariate normal distribution with $E(Z)=\mathbf{0}$ and $E(Z^2)=\boldsymbol{I}$, and let A be a nonnegative definite matrix. The distribution of $Z^{\mathrm{T}}AZ$ is the same as

$$\sum_{i=1}^{l}\lambda_i\chi_i^2(1)$$

where $\chi_1^2(1),\dots,\chi_l^2(1)$ are independent chi squared variables with 1 degree of freedom, and λ_1,...,λ_l are the eigenvalues of A.

Given Proposition (1) we need only show that A is has exactly q eigenvalues equal to 1. First we have

$$\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}\left\{E\left(\frac{\partial g}{\partial\boldsymbol{\theta}}\right)W_{\boldsymbol{\theta}}^{-1}E\left(\frac{\partial g}{\partial\boldsymbol{\theta}}\right)^{\mathrm{T}}-E\left(\frac{\partial g}{\partial\boldsymbol{\nu}}\right)W_{\boldsymbol{\nu}}^{-1}E\left(\frac{\partial g}{\partial\boldsymbol{\nu}}\right)^{\mathrm{T}}\right\}\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}$$

is nonnegative definite, idempotent and symmetric (see Qin and Lawless, 1994; Rao, 1973) with rank $p-(p-q)=q$ resulting in q eigenvalues equal to 1. The matrix A can be rewritten as

$$J^{-1}\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}\left\{E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)W_{\boldsymbol{\theta}}^{-1}E\left(\frac{\partial g}{\partial \boldsymbol{\theta}}\right)^{\mathrm{T}}-E\left(\frac{\partial g}{\partial \boldsymbol{\nu}}\right)W_{\boldsymbol{\nu}}^{-1}E\left(\frac{\partial g}{\partial \boldsymbol{\nu}}\right)^{\mathrm{T}}\right\}\left\{E\left(gg^{\mathrm{T}}\right)\right\}^{-1/2}\otimes 1_J 1_J^T.$$

The matrix $1_J 1_J^T$ has a single eigenvalue of J. Since the eigenvalues of A are equal to the product of the eigenvalues of each component we have the desired result. □